[

White Paper

Dell | Microsoft
Business Intelligence and Data Warehousing Reference Configuration – Performance Results Phase III

Performance of Microsoft® SQL Server™ 2008 BI and D/W Solutions on Dell™ PowerEdge™ Servers and Dell™ PowerVault™ Storage

Abstract

This white paper documents the performance of Microsoft SQL Server 2008 Business Intelligence and Data Warehouse solutions deployed on Dell PowerEdge servers with Dell PowerVault storage. Using the knowledge gained through joint development, testing and support with Microsoft, this white paper illustrates the cost-effective performance of Dell Solutions for SQL Server 2008.

November, 2008

November 2008 Rev. A02

Table of Contents

Introduction

Dell™ PowerEdge™ servers and Dell PowerVault™ storage systems are ideal choices to deploy highly available and enterprise mission critical Microsoft® SQL Server 2008™ Data Warehouses for Business Intelligence solutions.

Dell and Microsoft have worked together to run various tests and analyze performance results for a sample Business Intelligence (BI) and Data Warehouse (DW) solution. This performance testing was designed to validate the Dell PowerEdge server and PowerVault storage platforms for enterprise class BI and DW, to demonstrate the stability and performance of the systems, and to create Reference Configurations that can be used to assist customers in creating solutions for Business Intelligence and Data Warehousing. This is the third series of BI/DW reference configuration testing preformed by Dell and Microsoft. Phase I and II results can be found on line at www.dell.com/sqlbi. Phase III testing was preformed using the production release of SQL Server 2008 Enterprise Edition running on Windows Server 2008 Enterprise Editions. The previous tests we performed on dual-core AMD CPU's running SQL Server 2005 Enterprise Edition and Windows Server 2003 Enterprise Edition.

The focus of this white paper is to document the results of performance testing for the Dell/Microsoft Business Intelligence and Data Warehousing reference configurations. Using a simulated database of 10 terabytes and 20 Terabytes (using SQL Server 2008 page compression), tests were performed, including I/O stress, index creation of a table of five billion rows, concurrent query loading and backup and recovery timings. The methodology and results of these tests appear below.

Overview of Business Intelligence and Data Warehousing

In recent years the demands for Business Intelligence (BI) and Data Warehousing (DW) solutions have grown as new hardware and software technologies lowered cost and simplified implementation. As demand for BI and DW has grown, so too has the size and complexity of the databases grown. A unique set of tools and processes has been developed by Microsoft in order to meet the demands for this type of information management. The test methodology and corresponding results described in this paper illustrate that Dell servers and storage, teamed with Microsoft Windows Server and SQL Server 2008 are an excellent choice for BI/DW solutions with databases sizes into multi terabytes.

Business Intelligence

Business Intelligence (BI) is a broad term that refers to applications and technologies that are used to manage and analyze business operations data. BI is a way to monitor and report on various aspects of the business, from sales and marketing efforts, productivity and operations, to profitability. BI is designed to provide business decision makers with

information that enables optimal business decisions. The choice of what data is analyzed within a BI system is purely a business decision which is made possible by technology.

Some examples of how BI is used include: to inform management of the overall status and performance of the company, to provide information needed to be more competitive, to identify what areas of the business are in a positive or critical state, and to make decisions based on changes in the market, customer purchasing trends, product pricing and sales. BI allows analysts and managers to quickly understand the state of the business and make informed "real-time" decisions on how to adjust. In addition, tools like business score cards and executive dashboards provide instantaneous "state of the business" information to executive management. For these systems to be considered effective, large amounts of data must be analyzed very quickly.

BI is usually thought to operate against a data warehouse, but it is also possible to run BI against an online relational database. This white paper addresses BI primarily as it relates to a data warehousing. A data warehouse can be thought of as the corporation's repository of historical data. The data warehouse is typically very large in size and contains years of data that is typically used for analysis and reporting. These reports are then used in order to make decisions on direction for the company.

The technology of BI is made up of three main categories: the design of the data warehouse, the movement of data into the data warehouse, and reporting from the data. The design of the data warehouse includes either a relational database (SQL Server) or OLAP cubes (Analysis Services). Movement of data into the data warehouse can be done using any number of ETL (Extract Transform Load) tools, including SQL Server Integration Services (SSIS). Reporting and data mining can be accomplished with SQL Server Reporting Services (SSRS), Microsoft Excel, Microsoft Excel Services, or Performance Point Server 2007.

Data Warehousing

A data warehouse is a repository of historical information that is used to make organizational decisions. A data warehouse can contain data from many sources and can be very large in size, such as 10's or 100's of terabytes. How the data is stored depends in part on the software used to access that data. The performance of the data warehouse is dictated by both the software, hardware, and the configuration of each.

Microsoft solutions enable organizations to use several models to store data in their data warehouse. The model chosen depends on the data and how it will be accessed. Data in the data warehouse is stored in either a Relational Database Management System (RDBMS) such as SQL Server, in Analysis Services (or OLAP) cubes, or a combination. An Analysis Services cube can be thought of as a multidimensional abstraction of relational data.

Testing Dell Solutions for Microsoft SQL Server BI/DW

The BI/DW testing conducted for this white paper utilized the Microsoft Windows update database, Microsoft SQL Server 2008 Enterprise x64 Edition, Microsoft Windows Server 2008 Enterprise x64 Edition, Dell PowerEdge R905, 6950 and 2970 servers, and Dell PowerVault MD1000 storage systems. Windows Server 2008 Enterprise Edition x64 and SQL Server 2008 were installed on a Dell PowerEdge R905 quad-socket, quad-core server using AMD Opteron "Shanghai" CPU's with 64 GB of main memory and a 10 TB and a compressed 20 TB data warehouse was created. This data warehouse held the Microsoft Windows Update database, a Microsoft-provided data source simulating world-wide customer updates of their Microsoft products. Certain tables in this database have several billion rows of data.

A number of performance tests were run, including multiple user reporting, index creation, cube creation and data load. These tests exercised the I/O subsystem, memory subsystem and all 16 CPU cores in the system.

Overview of Microsoft BI Tools

Microsoft provides a broad range of tools for SQL Server Business Intelligence. Among these are SQL Server Reporting Services (SSRS), Excel 2007, and Performance Point Server 2007. The testing performed for this white paper focused on Microsoft SQL Server 2008, SSRS, Office Performance Point Server 2007 and Excel 2007.

Microsoft SQL Server Reporting Services

Microsoft SQL Server Reporting Services (SSRS) is a server based reporting tool designed to create, deploy and serve web based reports. These reports can be based on both relational and multidimensional databases. The SQL Server reporting architecture consists of the Report Server engine that runs the reports and the Report Server website which is used to display the reports.

One major advantage of SSRS is the wide range of data sources that can be used. This flexibility allows reports to be created from many different data sources in a heterogeneous environment. Currently SSRS can create reports from the following data sources:

- SQL Server Database
- SQL Server Analysis Services Cubes
- OLE DB
- Oracle database
- ODBC
- XML
- SAP NetWeaver BI
- Hyperion Essbase
- Teradata

Keep in mind that the report creation process consists of both the database queries and rendering of the report. Reports based on large amounts of data can consume significant CPU and memory resources on the report server; a best practice recommended by Dell is to deploy the report server engine from the database server on a separate server.

Microsoft Excel 2007

As communicated by Microsoft, Excel 2007 is evolving to become a primary end user reporting platform for business intelligence. Excel has been designed to integrate with web servers and other reporting tools in order to support "digital dashboards" for BI corporate performance management. Digital Dashboards are the future for monitoring key business indicators, but do not provide the full analytics of ProClarity or the reporting features of SSRS. In addition, Microsoft Excel Services in conjunction with Microsoft Office Sharepoint Server 2007 provides scalable and thin client rendering of Excel workbooks and charts that are published to the Sharepoint Server.

Microsoft Performance Point Server 2007

Microsoft Office PerformancePoint Server 2007 is an integrated performance management application designed to help improve operational and financial performance across all departments and all levels of your organization. With Office PerformancePoint Server 2007, you can monitor progress, analyze what is driving variances, and plan your business from budgeting to creating management reports. You can have metrics, key performance indicators (KPIs), and reports delivered to every desktop through intuitive scorecards, dashboards, and the easy-to-use 2007 Microsoft Office system environment.

Test Platform Configuration

The test configuration consisted of three tiers - a database server, an application server, and a reporting server. The following sections provide details of the hardware specifications and database configuration.

Dell Hardware

The database server consisted of a Dell PowerEdge R905 server with four, quad core AMD Opteron "Shanghai" CPU's running at 2.59 GHz, 64 GB of RAM and 336 SAS disk drives for data and 60 drives for backups. (The Dell PowerEdge R905 includes seven PCI Express slots, 2 x8 slots and 5 x4 slots)

The 336, 3.5" disk drives resided in 24 - Dell PowerVault MD1000 storage cabinets and were controlled by 6 Dell Perc 6/E controllers from the Dell R905 server. Each MD1000 has one extra disk for hot spares in the case of disk failure. On the R905, 1 additional Perc 6/E controller was used for the four additional MD1000's used for backup space and TempDB. A summary of the hardware and software for the database server is listed in the table below.

Dell PowerEdge R905 Database Server	
• Dell PowerEdge R905 • 64 GB RAM • 4 x Four Core AMD Opteron "Shanghai" CPU's, 2.59 Ghz • 4 x 73 GB, 15K SAS Internal Disks • Seven external SAS Disk Controllers; Dell Perc 6/E. 6 controllers are used for the user databases and the other 1 for backups, Tempdb and temporary work space. • 24 Dell PowerVault MD1000 Storage Arrays • 308 x 73 GB 15K SAS Drives and 28 x 146 GB 15K SAS Drives (total = 332)Windows Server2008 Enterprise x64 Edition • SQL Server 2008 Enterprise Edition x64	• Dell Server/Storage Management software • Microsoft Operations Manager (MOM) software • Four additional Dell PowerVault MD1000 Storage Arrays for backup • 60 x 73 GB 10K SAS Drives for backup files and source data.

Data Type	Storage	# Disks	RAID
SQL Transaction Logs	Internal	4	2 x RAID 1 pairs
SQL User DB Data Files	External	336	28 x RAID 10 pairs
SQL Tempdb data and log	External	8	4 x RAID 1 pairs
Database Backups	External	52	RAID 0

For the database server storage, the user database files and the SQL Server tempdb files were physically separated in order to achieve optimal performance. The 336 external disks were configured as 24 RAID 10 pairs and one hot-spare per MD1000 array (24 hot spares total). The SQL Server transaction logs were placed on the internal disks.

In addition, disk storage for database backups was configured across 52 x 73 GB 10K SAS drives as four RAID 0 sets (on 4 separate MD1000's). These drives were configured as RAID 0 (striped, with no fault tolerance) for testing purposes only, as RAID 0 is not recommended for production deployments.

Hardware and software configurations for the Analysis and Reporting servers are in the following tables:

Analysis Server	
• Dell PowerEdge 6950 • 32 GB RAM • 4 x Dual Core AMD Opteron CPUs • 4 x 146 GB, 15K SAS Internal Disks • Two external SAS Disk Controllers • Two Dell PowerVault MD1000 Storage Array • 30 x 73 GB 10K SAS Drives	• Microsoft Windows 2008 Enterprise x64 Edition • Microsoft SQL Server 2008 Enterprise Edition x64 (per Processor licensing) • Dell Server/Storage Management Software • MOM

Reporting Server	
• Dell PowerEdge 2970 • 32 GB RAM • 2 x Dual Core AMD Opteron CPUs • 4 x 73 GB, 10K SAS Internal Disks	• Microsoft Windows 2008 Enterprise x64 Edition • Microsoft SQL Server 2008 Enterprise Edition x64 (per Processor licensing) • Dell Server/Storage Management Software • MOM

PCI Express

The I/O performance achieved in testing was made possible by the use of the PCI Express (PCIe) bus[1]. The PCIe slots are significantly faster than PCI-X slots as seen in the table below.

Bus and Frequency	Peak Transfer Rate (64-bit)
33 MHz PCI	266 Mbytes/sec
66 MHz PCI	532 Mbytes/sec
100 MHz PCI-X	800 Mbytes/sec
133 MHz PCI-X	1 Gbytes/sec
PCIe x1	5 Gbps
PCIe x4	20 Gbps
PCIe x8	40 Gbps
PCIe x16	80 Gbps

Figure 1 - PCIe Performance

[1] Dell Technology White Paper. PCI Express Technology. February 2004. Jim Brewer, Joe Sekel.

Storage

Serial Attached SCSI or SAS drives provide the flexibility of SCSI disks with high throughput. There are two factors that affect the performance of your I/O subsystem; the throughput or speed of the bus, and the speed of the drive (i.e. rotational speed, seek time). The Serial Attached SCSI (SAS) and traditional SCSI use the same underlying disk technology but the electronics that control the bus is different. Currently SAS bandwidth is 3 GB/sec, with 6 GB/sec and 12GB/sec planned for the future. This table compares SAS, Ultra-320 SCSI, and SATA drives.

	SAS	SCSI – U320	SATA	Notes
Bus Bandwidth	3 GB/sec	320 MB/sec	3 Gbits/sec	Note Gigabit/sec on SATA vs. Gigabytes/sec
Rotational Speed	15,000 RPM	15,000 RPM	7200 RPM	
Avg. Seek Time	3.5 ms	3.5 ms	20 ms	

Figure 2 - SCSI and SAS Performance

The results of disk testing with SAS arrays are found in the section “Storage Testing” later in this paper.

Database

The database used for testing was based on the Microsoft Worldwide Windows update database and was created at approximately 1.8 TB in size. Additional copies of the large fact tables were made to bring the data size up to 10 TB plus in size and even 20 TB in size with page compression in SQL Server 2008. The databases on the system and file sizes are listed in the following table, with WinUpdtNew being the data warehouse:

DB Name	DB Size (MB)	Log Size (MB)	Log Space Used (%)
master	4.44	0.49	61.11
tempdb	8	0.49	77.38
model	1.19	0.49	69.84
msdb	5.44	0.49	61.9
WinUpdtNew	11,885,044	236,000	0.71

Figure 3 - Windows Update Database File Sizes

Row counts on some of the larger tables exceed one billion rows as shown here:

Name	**Row Count**
Fact**********	5,011,894,223
Fact**********	3,627,211,650
Fact**********	2,234,928,020
Fact**********	177,702,135
Fact**********	100,000,000
Dim********	77,080,636
Fact**********	1,508,281

Figure 4 - Windows Update Database Table Row Counts

Note: The table names have been altered for privacy purposes.

Performance Results

A number of batch and online tests were performed to stress I/O performance and overall SQL Server 2008 performance. Full system testing was done using OLTP multi-user loads as well as with index and OLAP cube creation.

Storage Testing

Before starting the Microsoft SQL Server 2008 testing, a number of tests were run using the simulation program sqlio. The sqlio program is designed to simulate SQL Server I/O operations and provides detailed information on the response times for I/Os. By using sqlio it was possible to completely stress the I/O subsystem without having to worry about SQL Server contention and CPU limitations. In addition, sqlio allows specification of block sizes, outstanding I/Os, and other test options that provide a range of results based on the input parameters.

R905 Test Results

A read only sqlio test was run against all 24 user disk volumes (i.e. all 336 disks).
With this test, a very fast 71,044 IOPS were achieved with a 4.44 GB/sec read throughput at a low latency of 8 ms, as shown in the highlighted row of Figure 5. This is a significant throughput achievement and far greater than the 2.7 GB/sec read performance achieved on Dell servers in December of 2007.

Block Size	Outstanding IO's	Read or Write	IOPS	MB/Sec	Latency(ms)
64	4	Read	67,113	4,194	2
64	8	Read	70,738	4,421	5
64	**12**	**Read**	**71,044**	**4,440**	**8**
64	16	Read	67,745	4,234	12
64	20	Read	67,904	4,244	15
64	32	Read	66,649	4,165	25

Figure 5 - SQLIO Test Results for 24 MD1000 Disk Cabinets (336 total disks)

Figure 6 shows a graphical view of the outstanding I/Os and throughput from this test.

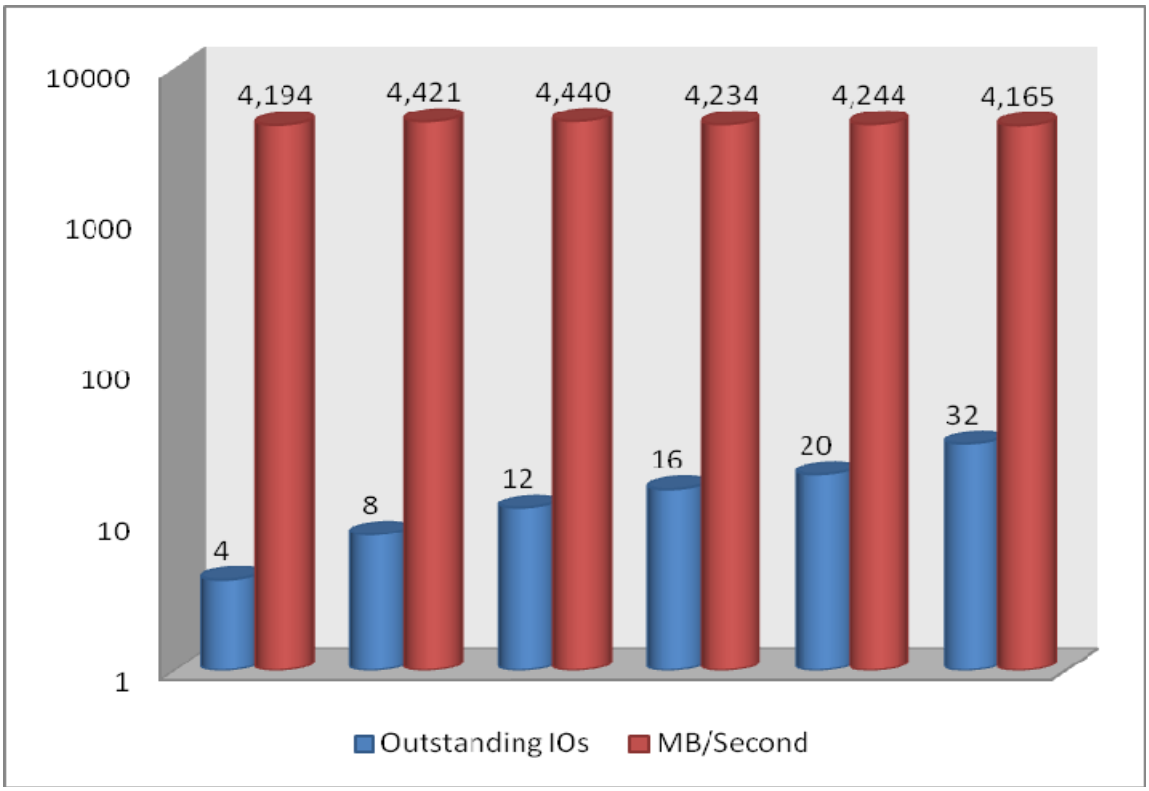

Figure 6 - Outstanding I/Os vs. Throughput with 5 Disk Cabinets

Index Creation

In order to characterize the performance of the index creation operation, a non-clustered index was created and timed on a table containing over 5 billion rows. This index creation took 11 minutes and 10 seconds. The previous best time on the largest Dell servers in our phase I and II testing completed in 2007 was 53 minutes 26 seconds using SQL Server 2005 database. This type of performance makes it possible to maintain very large databases within normal maintenance windows. This index build time of 11 minutes and 13 seconds was accomplished using 16 cores (4 socket x 4 core) on the R905 server and a compressed SQL Server 2008 database which has 20 TB's of user data.

As underlying table data is updated, indexes can become fragmented. This is especially true of clustered indexes. As indexes become fragmented, index lookup performance can become significantly degraded. Thus indexes must be rebuilt on a regular basis. The time it takes to rebuild indexes can be significant, especially on large tables. With the Dell reference platform, we were able to rebuild large indexes very quickly.

Large Table Queries

Two typical data warehouse queries were created to sum measures across a 5.1 billion row fact table and a 2.35 billion row fact table.

Query	Dell R905 Server (AMD Based)
Select Sum of 3 measures on a 5.1 billion row fact tables with no indexes.	4 minutes 31 Seconds
Select Sum of 3 measures and group by 4 columns on 2.35 billion row fact table	2 minutes 14 Seconds

Note: With SQL Server 2005 and 2008 index maintenance can be performed as an online operation, allowing users to continue accessing the underlying table data.

Concurrent Query Loads

By simulating hundreds of users, many components of the system were tested. These tests stressed I/O, memory and CPU and provided both throughput and response time results. By performing an end-to-end test, the entire system is stressed, better reflecting what the end user might experience.

This user simulation test was taken from Microsoft Project REAL, which can be found at the web site http://www.microsoft.com/sql/solutions/bi/projectreal.mspx. This project was designed to create best practices developed by real customer scenarios, with Microsoft Visual Studio Team Edition used to generate user loads.

The test system was able to successfully sustain 250 concurrent users with response times under 5 seconds, as shown in Figure 8. (User response time is represented by the blue line and user load, the red line).

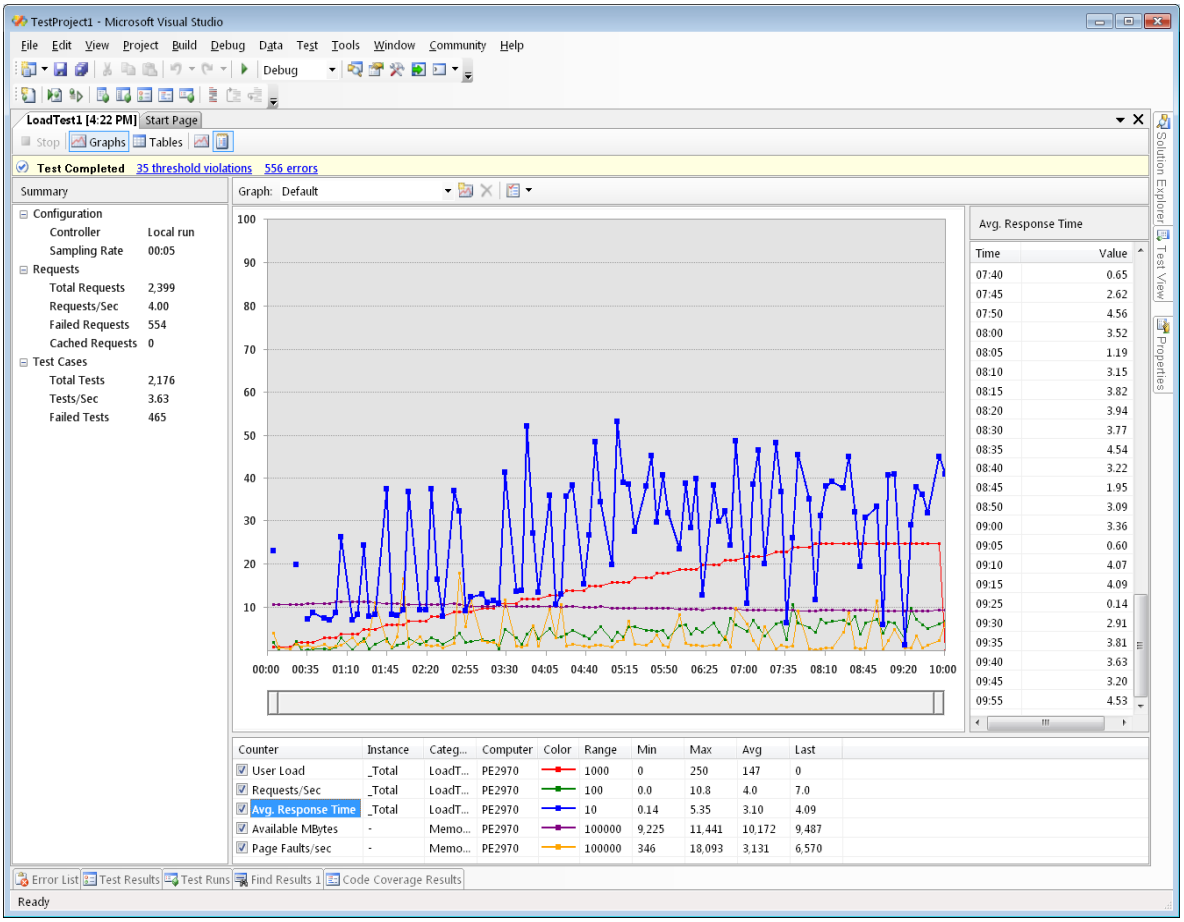

Figure 7 - Concurrent Query Performance Results

Conclusions

Dell Solutions for SQL Server 2008 are designed to simplify operations, improve utilization and cost-effectively scale as your needs grow over time. This performance white paper illustrates the efficiencies of deploying SQL Server 2008 Business Intelligence and Data Warehousing solutions on Dell PowerEdge servers and Dell PowerVault MD1000 storage arrays.

The results described here are intended to provide a performance baseline for Dell Solutions for SQL Server 2008 BI/DW. To learn more about deploying SQL Server 2008 on PowerEdge server and Dell storage, please visit www.dell.com/sql or contact your Dell representative for up to date information on Dell servers, storage and services for SQL Server 2008 solutions.

Index of Figures